

Next-Generation Backup and Recovery for MongoDB

Rubrik NoSQL Protection provides a cloud-scale, application-centric data management platform enabling organizations to protect, mobilize, and monetize all of their MongoDB application data across private cloud, hybrid cloud, and public cloud environments.

KEY BENEFITS

Application-Centric

- Application and cluster-consistent point-in-time backups
- Scalable versioning supports large clusters
- Flexible deployment with API-based architecture and native-UI interface

Recover in Minutes, Not Hours

- Single-click, fully orchestrated recovery
- Recovery to same or different-sized clusters
- Granular recovery for fastest RTO

Increase Ops Efficiency

- Semantic deduplication cuts backup storage requirements up to 70%
- Automate Test/Dev refresh
- Use backup copies for migrations, Test/Dev, and database cloning

THE CHALLENGE

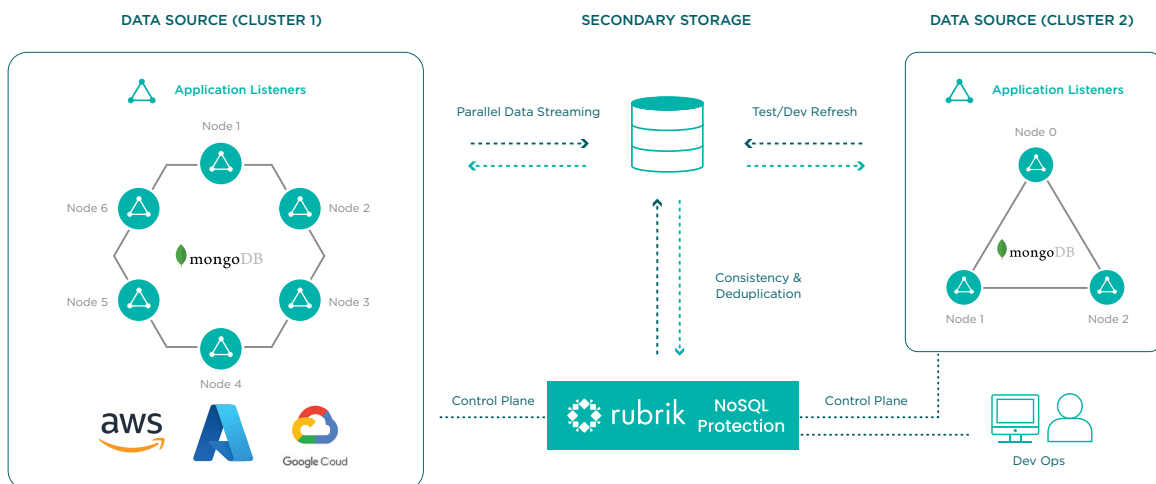
Businesses are in the midst of a digital transformation journey. According to research from IDC, 70% of CIOs have a cloud-first strategy. They want to harness the power of the cloud to drive growth by delivering new customer-centric products and services while also driving greater operational efficiency. To handle the data requirements of these modern high-volume, high-ingestion-rate, and real-time applications, enterprises are turning to scalable, nonrelational databases such as MongoDB rather than traditional scale-up database and storage approaches.

However, this fundamental shift raises critical issues in the life cycle of data management and data protection. Traditional backup and recovery products were originally designed for small-scale databases, tape-based storage media, and legacy on-premises architectures. This leaves modern applications built on nonrelational databases and Big Data filesystems exposed to data loss and downtime.

THE SOLUTION: RUBRIK NOSQL PROTECTION

Rubrik NoSQL Protection is the industry's first and only scale-out data protection software solution to deliver scalable and reliable backup and recovery for modern applications built on MongoDB databases. With NoSQL Protection, enterprises can deploy business-critical applications on MongoDB and be confident in the recoverability of data and the ability to maintain high application uptime.

RUBRIK NOSQL PROTECTION FOR MONGODB: REFERENCE ARCHITECTURE



FEATURES AND BENEFITS

Rubrik NoSQL Protection is built to address the data protection needs of modern, cloud-native applications deployed on MongoDB. Unique capabilities include:

Scale-Out Architecture

Rubrik NoSQL Protection is founded upon *Consistent Orchestrated Distributed Recovery™* (CODR™), NoSQL Protection's cloud-first, scale-out data management architecture that enables customers to meet their data protection requirements for MongoDB. CODR uses elastic-compute services that can be autoscaled with load, removing the dependency on media servers. CODR also transfers data in parallel to and from file-based and object-based secondary storage for multiple use cases, including data protection and Test/Dev refresh. To simplify the data recovery process and to avoid vendor lock-in, protected data is stored in the database-native format.

Continuous Backup

By using native application intelligence, NoSQL Protection creates point-in-time, consistent backup copies of MongoDB collections (both sharded and unsharded) at user-specified intervals—a concept that we refer to as *cluster-consistent versioning*. NoSQL Protection can produce these cluster-consistent versions across all shards without quiescing the MongoDB database. You can generate backups at a user-specified time interval and at any granularity (collection-level or entire database), providing operational ease of use to database administrators. And with NoSQL Protection, backup operations are resilient to failovers (primary switch) and failures (node).

Fully Orchestrated and Granular Recovery

NoSQL Protection provides fully orchestrated, any-point-in-time recovery. You can select granular recovery based upon either time or query for optimal recovery time objective (RTO)/recovery point objective (RPO) and to support governance requirements like

the European Union General Data Protection Regulation (GDPR). You can recover MongoDB collections directly back into the same MongoDB database (operational recovery). You also can recover them to a different MongoDB database instance (e.g., Test/Dev refresh) with a different topology for which the number of nodes on the destination cluster differs from the node count of the source cluster.

Rubrik NoSQL Protection supports all combinations of recovery—sharded to unsharded, unsharded to sharded, sharded to sharded—thus reducing the operational burden of refreshing Test/Dev clusters for Continuous Development DevOps environments. Further, the recovery process deals with only the logical data, making it three times faster than traditional approaches. During recovery, the data is directly transferred from secondary storage into target databases, removing intermediary media servers, resulting in the lowest possible RTO.

Semantic Deduplication

NoSQL Protection includes *semantic de-duplication*, an industry-first capability that reduces the cost of storing backups of distributed databases over their retention period. These space-efficient backups dramatically reduce the overall storage footprint resulting in up to a 70% reduction in backup storage requirements.

Infrastructure and Storage Independence

NoSQL Protection is elastic-compute software that can be deployed on a physical server, a virtual machine, or any cloud compute instance (e.g., Amazon Elastic Compute Cloud [Amazon EC2], Microsoft Azure, Google Cloud Platform and Oracle Cloud). You can backup data to any Network File System (NFS) or object storage on-premises or in a public cloud (e.g., Amazon Simple Storage Service [Amazon S3]). In addition to command-line interfaces and RESTful APIs, you can use the NoSQL Protection consumer-grade UI to manage your data protection environment.

RUBRIK NOSQL PROTECTION FOR MONGODB: COMPATIBILITY MATRIX

MongoDB database version support	MongoDB 3.0, 3.2, 3.4, 4.x			
Deployment	On-premises	AWS	Google Cloud	Microsoft Azure
Secondary storage support	NFS, Object Storage	Amazon S3	Google Cloud Storage	Azure Blob Storage
Rubrik NoSQL Protection Software Nodes (3-nodes)	RHEL/Centos 6.x			
	16-core, 32 GB memory	EC2 M4.2xlarge or above	Standard 8 vCPU, 30 GB RAM	Azure Blob Storage
	256 GB local storage (SSD)			